A New Algorithm for Solving Ring-LPN with a Reducible Polynomial

Qian Guo, Thomas Johansson, *Member, IEEE,* and Carl Löndahl

arXiv:1409.0472v1 [cs.CR] 1 Sep 2014

***Abstract*—The LPN (Learning Parity with Noise) problem has recently proved to be of great importance in cryptology. A special and very useful case is the RING-LPN problem, which typically provides improved efficiency in the constructed cryptographic primitive. We present a new algorithm for solving the RING-LPN problem in the case when the polynomial used is reducible. It greatly outperforms previous algorithms for solving this problem. Using the algorithm, we can break the Lapin authentication protocol for the proposed instance using a reducible polynomial, in about 2^{70} bit operations.**

***Index Terms*—Birthday attacks, Fast Walsh-Hadamard Transform, Lapin, LPN, RING-LPN.**

I. INTRODUCTION

LIGHT-WEIGHT cryptography is a field of cryptography inclined towards efficient cryptographic implementations, as a response to the demands when using highly constrained hardware in low-cost devices, such as passive RFID-tags and smart cards.

There are trade-offs to consider, e.g., security, memory and performance. Different constructions appear in different ends in the trade-offs; for instance, AES and stream ciphers can be implemented efficiently in hardware but do not offer *provable security*. Quite recently, a new trend arose in this area, building cryptographic primitives from problems in learning theory. Problems based on learning theory provide a complexity theoretical foundation, on which the security of the cryptosystem can be based upon. They also have the property of being easy and efficiently implemented, thereby making them appealing in light-weight cryptography.

A. The LPN Problem

Being a central problem in learning theory, the LPN problem (*Learning Parity with Noise*) has shown to be of significance in the field of cryptography. It is a supposedly hard problem[1], and is not known to be susceptible to quantum attacks, unlike some other classically hard problems such as factoring and the discrete log problem.

The problem can briefly be described as follows. Let $\mathbf{s}$ be a k-dimensional binary vector. We receive a number N of noisy versions of scalar products of $\mathbf{s}$ from an oracle Π_{LPN}, and our task is to recover $\mathbf{s}$.

Let $\mathbf{y}$ be a vector of length N and let $y_i = \langle \mathbf{s}, \mathbf{r}_i \rangle$. For known random k bit vectors $\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_N$, we can easily reconstruct an unknown $\mathbf{x}$ from $\mathbf{y}$ using linear algebra. In the LPN problem, however, we receive instead noisy versions of $y_i, 1 \leq i \leq N$.

Writing the noise in position i as $e_i,\ i = 1, 2, \ldots, N$, and assuming each $\mathbf{Pr}\left[e_i = 1\right]$ to be small, we obtain

$$v_i = y_i + e_i = \langle \mathbf{x}, \mathbf{r}_i \rangle + e_i.$$

In matrix form, the same is written as $\mathbf{v} = A\mathbf{s} + \mathbf{e}$, where $\mathbf{v} = \begin{bmatrix} v_1 & v_2 & \cdots & v_N \end{bmatrix}$, and the matrix A is formed as $A = \begin{bmatrix} \mathbf{r}_1 & \mathbf{r}_2 & \cdots & \mathbf{r}_N \end{bmatrix}$.

B. Constructions and Variations of LPN

The LPN problem and its variations have been employed as the underlying hard problem in a wide range of public-key cryptosystems, identification and authentication protocols, and zero-knowledge proofs.

The first actual usage of the LPN problem in cryptographic context can be traced back to 2001, when the Hopper-Blum (HB) identification protocol [13] was proposed. Being an intentionally minimalistic LPN based protocol, it was designed so that it could be executed by humans using only pen and paper. Much due to its simplicity, it is secure only in the *passive* attack model. A couple of years later, Juels and Weis [14] along with Katz and Shin [15] proposed a modified scheme, extending HB with one extra round. The modified scheme was named HB$^+$. Contrary to its predecessor, HB$^+$ was designed to be also secure in the *active* attack model. However, it was discovered by Gilbert et al. [11] that the HB$^+$ protocol is susceptible to *man-in-the-middle attacks* disproving active attack model security. The same authors [10] proposed later on another variation of the Hopper-Blum protocol called HB$^\#$, designed to resist their previous attack [11]. Apart from repairing the protocol, they solved the long-lived issue with large key-size or communication complexity by introducing the use of a slight variation of LPN, called TOEPLITZ-LPN. Although the use of TOEPLITZ-LPN has no documented weaknesses, its hardness remains unknown as of today.

During the time from when Hopper and Blum pioneered the use of LPN until today, a plethora of different proposals have hit the cryptographic society. Some of the most important ones are the proposals by Klitz et al. [18] and Dodis et al. [7] that showed how to construct message authentication codes based on LPN. The existence of MACs allows one to construct identification schemes that are provably secure against active attacks.

Most recently Heyse et al. [12] proposed a two-round identification protocol called Lapin. The protocol is based on RING-LPN rather than LPN. Using the inherent properties of

The authors are with the Department of Electrical and Information Technology, Lund University, Box 118, SE-22100 Lund, Sweden (e-mail: qian.guo@eit.lth.se; thomas.johansson@eit.lth.se; carl@grocid.net).

[1]LPN with adversarial errors is $\mathcal{NP}$-hard.

TABLE I
COMPARISON OF DIFFERENT ALGORITHMS FOR ATTACKING LAPIN WITH REDUCIBLE POLYNOMIAL.

Algorithm	Complexity ($\log_2$)		
	Queries	Time	Memory
Levieil-Fouque [20]	82.0	103.4	100.6
Bernstein-Lange [4]	79.3	102.9	97.9
Our attack (search)	63	71.9	70.0
Our attack (decision)	62	70.0	69.0

rings, the proposed protocol becomes very efficient and well-suited for use in constrained environments. Briefly, in RING-LPN, the oracle returns instead elements v, $v = s \cdot r + e$, from a polynomial ring $\mathbb{F}_2[x]/(f)$, i.e., $v, s, r, e \in \mathbb{F}_2[x]/(f)$. The problem can use either an irreducible polynomial f, of a reducible one. The choice of a reducible polynomial can make use of the Chinese Remainder Theorem (CRT) to provide a very efficient implementation of the cryptographic primitive.

C. Attacks

1) Attacks on Standard LPN: The LPN problem can be viewed as a general decoding problem in coding theory. However, the usual choice of parameters for the LPN problem is deviating from standard parameters for a decoding problem. The LPN problem typically allows a very large amount of oracle queries, say $N = 2^{60}$ or even much larger; therefore, standard algorithms for the general decoding problem, eg. information-set decoding, will not always be very efficient for such cases.

Instead, a slightly different type of algorithms have been suggested, among which we find the BKW algorithm [2] proposed by Blum et al. It was later refined by Levieil and Fouque [20]. They gave better attacks by employing the Fast Walsh-Hadamard Transformation technique, and estimated the security level of LPN problems with different parameters. Fossorier et al. [8] suggested an algorithm which further improved the complexity by utilizing techniques from the area of fast correlation attacks. In the recent paper by Kirchner [19], it is proposed to exhaust the search space of the error, rather than the state. For a certain class of instances, i.e., when the error rate is low, Kirchner's technique greatly improves attack complexity.

Lyubashevsky's work [21] moved forward in the other research direction, i.e., when the number of required samples is bounded. Compared with BKW algorithm, he presented an asymptotically efficient algorithm with slightly increased time complexity, i.e., from $2^{\mathcal{O}(n/\log n)}$ to $2^{\mathcal{O}(n/\log\log n)}$, while obtaining much query efficiency.

2) Attacks on Ring-LPN: As the RING-LPN instances are standard LPN instances, the attacking algorithm for the latter one is applicable to its ring instance. The pioneering researchers (eg. Lapin [12]) used the hardness level of the LPN problem obtained from [20] to measure the security of their authentication protocol based on the RING-LPN problem. Almost at the same time, Bernstein and Lange [4] realized that simply ignoring the ring structure is inappropriate since this special algebraic property may reveal information about the secret, and subsequently derived an improved attack taking advantage of both the ring structure and Kirchner's technique. Their attack is generic since it applies to RING-LPN implemented with both reducible and irreducible polynomials, and is advantageous in the sense of memory costs as well as query complexity. However, even for the time-optimized case[2], with around 2^{81} bits of memory, it requires quite a large number of bit operations, i.e., about 2^{88}, far away from breaking the 80-bit security of Lapin protocol.

D. Our Contribution

We propose a new algorithm to solve the reducible case of RING-LPN. By investigating more on the properties of the ring structure and the reducibility of the polynomial, we demonstrate that if the minimum weight of the linear code defined by the CRT transform is low, then the problem is effortless to solve, hence providing a design criteria for cryptosystems based on the hardness of RING-LPN with a reducible polynomial.

We then specify the attack for Lapin [12] and obtain a complexity gain that makes it possible to break the claimed 80-bit security. In Table I, we compare the complexity of our algorithm with the best known algorithms[3] designed to solve LPN and RING-LPN. The time complexity is measured in bit operations and memory complexity is measured in bits.

The organization of the paper is as follows. In Section II, we give some preliminaries and introduce the RING-LPN problem in detail. We describe our new generic attack in Section III and then present a special version that is efficient for the proposed reducible instance of Lapin in Section IV. In Section V we analyze its complexity. The numerical results when the algorithm is applied on Lapin are given in Section VI and Section VII concludes the paper.

II. THE RING-LPN PROBLEM

A. Polynomials and Rings

Consider a polynomial $f(x)$ over $\mathbb{F}_2$ (simply denoted f). The degree of f is denoted by $\deg f$. For any two polynomials f, g in the quotient ring $\mathbb{F}_2[x]$, long division of polynomials tells us that there are unique polynomials q, r such that $g = qf + r$. The unique polynomial $r \in \mathbb{F}_2[x]$ with $\deg r < \deg f$ is the representative of g in the quotient ring $\mathbb{F}_2[x]/(f)$ and

[2]We found that with parameters $q = 2^{58.59}$, $a = 6$, $b = 65$, $l = 53$ and $W = 4$, the complexity is $2^{85.9}$. The time complexity is slightly lower than stated in [4].

[3]We choose parameters to optimize their time complexity.

is denoted $g \bmod f$. We define R to be the quotient ring $\mathbb{F}_2[x]/(f)$. So R consists of all polynomials in $\mathbb{F}_2[x]$ of degree less than $\deg f$ and arithmetics are done modulo f.

If the polynomial f factors such that every factor is of degree strictly less than f, then the polynomial f is said to be *reducible*; otherwise, it is said to be *irreducible*. When a reducible polynomial factors as $f = f_1 f_2 \cdots f_m$, where f_i is relatively prime to f_j for $1 \leq i, j \leq m$ and $i \neq j$, there exists a unique representation for every $r \in R$ according to the Chinese Remainder Theorem, i.e.,

$$r \mapsto (r \bmod f_1, r \bmod f_2, \ldots, r \bmod f_m).$$

B. Distributions

Let Ber_η be the Bernoulli distribution and X be a binary random variable. It is said that X is distributed according to the Bernoulli distribution with parameter η, if $\mathbf{Pr}[X = 1] = \eta$ and $\mathbf{Pr}[X = 0] = 1 - \mathbf{Pr}[X = 1] = 1 - \eta$. Then, the bias ϵ of X is given by $\mathbf{Pr}[X = 0] = \frac{1}{2}(1 + \epsilon)$, i.e. $\epsilon = 1 - 2\eta$. Let $r \xleftarrow{\$} \mathsf{Ber}_\eta^R$ denote that the coefficients of the ring element $r \in R$ are drawn randomly according to the distribution Ber_η. The uniform distribution is denoted U. Whenever we draw an element uniformly from R, we denote it $r \xleftarrow{\$} \mathsf{U}^R$.

C. Formal Definition of Ring-LPN

Being a subclass of LPN, the RING-LPN problem is defined similarly. Fix an unknown value $s \in R$, where $s \xleftarrow{\$} \mathsf{U}^R$. We can request samples depending on s through an oracle, which we define as follows.

*Definition 1 (*RING-LPN *oracle):* A RING-LPN oracle $\Pi^{\epsilon}_{\text{RING-LPN}}$ for an unknown polynomial $s \in R$ with $\eta \in (0, \frac{1}{2})$ returns pairs of the form

$$(r,\ r \cdot s + e),$$

where $r \xleftarrow{\$} \mathsf{U}^R$ and $e \xleftarrow{\$} \mathsf{Ber}_\eta^R$. An extreme case is, when η is exactly $\frac{1}{2}$, that the oracle $\Pi^{0}_{\text{RING-LPN}}$ outputs a random sample distributed uniformly on $R \times R$.

The problem is now to recover the unknown value s after a number q of queries to the oracle. We define the search problem version of RING-LPN in the following way.

*Problem 1 (*RING-LPN*):* The search problem of RING-LPN is said to be (t, q, δ)-solvable if there exists an algorithm $\mathcal{A}(\Pi^{\epsilon}_{\text{RING-LPN}})$ that can find the unknown polynomial $s \in R$ in time at most t and using at most q oracles queries such that

$$\mathbf{Pr}[\mathcal{A}(\Pi^{\epsilon}_{\text{RING-LPN}}) = s] \geq \delta.$$

The decisional RING-LPN assumption, states that it is hard to distinguish uniformly random samples from pairs from the oracle $\Pi^{\epsilon}_{\text{RING-LPN}}$. It can be expressed as follows.

Problem 2 (Decisional RING-LPN*):* The decision problem of RING-LPN is said to be (t, q, δ)-solvable if there exists an algorithm $\mathcal{D}$ such that

$$\left|\mathbf{Pr}[\mathcal{D}(\Pi^{\epsilon}_{\text{RING-LPN}}) = \mathsf{yes}] - \mathbf{Pr}\left[\mathcal{D}(\Pi^{0}_{\text{RING-LPN}}) = \mathsf{yes}\right]\right| \geq \delta,$$

and $\mathcal{D}$ is running in time t and making q queries.

The hardness of RING-LPN is unknown, but the LPN problem has been shown to be $\mathcal{NP}$-hard in the worst-case. The assumption is that RING-LPN is also hard.

In the paper by Heyse et al. [12], it was proposed to use RING-LPN as the underlying hard problem to build an authentication protocol. The security relies on the assumption that RING-LPN is as hard as LPN. However, this is a conjecture as there is only a reduction from RING-LPN to LPN, but not the converse. In the following, we show how to reduce RING-LPN to LPN.

1) Transforming Ring-LPN to LPN: Given the polynomial r with $\deg r = t$, we denote $\mathbf{r}$ as its coefficient vector, i.e., if r equals $\sum_{i=0}^{t} r_i x^i$, then $\mathbf{r} = \begin{bmatrix} r_0 & r_1 & \cdots & r_t \end{bmatrix}$. With this notation, we can define a mapping from one RING-LPN instance to d standard LPN instances represented in the following matrix form:

$$\begin{array}{rccc} \tau: & R \times R & \to & \mathbb{F}_2^{t \times t} \times \mathbb{F}_2^t \\ & (r, r \cdot s + e) & \mapsto & (A, A\mathbf{s} + \mathbf{e}) \end{array} \quad (1)$$

where the i-th column of the matrix A is the transposed coefficient vector of $r \cdot x^i \bmod f$.

2) Ring-LPN with a Reducible f: In this paper we consider the RING-LPN problem with the ring $R = \mathbb{F}_2[x]/(f)$, where the polynomial f factors into distinct irreducible factors over $\mathbb{F}_2$. That is, the polynomial f is written as $f = f_1 f_2 \cdots f_m$, where each f_i is irreducible. One of the specified instances of Lapin [12] uses a product of five different irreducible polynomials. This instance is the main target for cryptanalysis in this paper.

D. Basics on Coding Theory

For later use, we end this section by reviewing some basics on *linear codes* over the binary field. A linear code $\mathcal{C}$ is a k-dimensional subspace of an n-dimensional vector space. The elements of $\mathcal{C}$ are called *codewords* and the Hamming weight of a codeword is defined as number of non-zero entries. The *minimum distance* of a linear code is defined as the lowest weight among the weights of all codewords.

Definition 2 (Generator matrix): A generator matrix G for $\mathcal{C}$ is defined as a $k \times n$ matrix whose span is the codeword space $\mathcal{C}$, i.e.,

$$\mathcal{C} = \left\{\mathbf{x}G : \mathbf{x} \in \mathbb{F}_2^k\right\}.$$

The existence of codes with specific parameters is sometimes guaranteed through the famous GV bound.

Theorem 1 (Gilbert-Varshamov bound): Let n, k and d be positive integers such that

$$\mathrm{Vol}(n-1, d-2) < 2^{n-k},$$

where $\mathrm{Vol}(n-1, d-2) = \sum_{i=0}^{d-2} \binom{n-1}{i}$. Then, there exists an $[n, k]$ linear code having minimum distance at least d.

A *random* linear code is a code $\mathcal{C}$, where the entries of the generator matrix G has been selected according to an i.i.d. uniform distribution. It is known that the minimum distance d for a random linear code asymptotically follows the GV bound.

III. The New Algorithm for Ring-LPN with Reducible Polynomial

The purpose of the paper is to describe a new algorithm (Algorithm 1) for the Ring-LPN problem with a reducible polynomial. We describe the algorithm as follows. First, we reduce the problem into a smaller one, while keeping the error at a reasonable level. Then, we further reduce the unknown variables using well-established collision techniques. The last step consists of exhausting the remaining unknown variables in an efficient way.

A. A Low-Weight Code from the CRT Map

Let $f = f_1 f_2 \cdots f_m$ be a reducible polynomial. In the following, the underlying irreducible polynomial f_i in the quotient ring $\mathbb{F}_2[x]/(f_i)$ is fixed. Therefore, w.l.o.g., we denote it as f_1 for notational simplicity. Let $\deg f = t$ and $\deg f_1 = l$.

Proposition 1: There exists a (surjective) linear map from $\psi : R \to \mathbb{F}_2[x]/(f_1)$ determined by the CRT transform. The linear map can be described as a $[t, l]$ linear code[4] with generator matrix G_ψ, in which the i-th column is the transposed coefficient vector of the polynomial $x^{i-1} \bmod f_1$.

More specifically, a received sample $(r, r \cdot s + e)$ from the Ring-LPN oracle $\Pi^{\epsilon}_{\text{Ring-LPN}}$ can be transformed into a considerably smaller instance

$$\psi : (r, r \cdot s + e) \mapsto (r \bmod f_1, (r \bmod f_1) \cdot (s \bmod f_1) + e \bmod f_1).$$

For simplicity, we write $\hat{r} = r \bmod f_1$, $\hat{s} = s \bmod f_1$ and $\hat{e} = e \bmod f_1$. As before, we may write $\hat{r} = \sum_{i=0}^{l-1} \hat{r}_i x^i$, etc.

The new instance has a smaller dimension, as $\deg f_1 < \deg f$. However, the distribution of the errors is also changed. The error distribution in the larger ring is $\mathsf{Ber}_{\frac{1}{2}(1-\epsilon)}$, but in the smaller ring each noise variable $(\hat{e} = e \bmod f_1)$ is a sum of several entries from e. The number of noise variables that constitutes a new error position $(\hat{e} = e \bmod f_1)$ depends entirely on the relation between f and f_1.

The following is an example that chooses f_1 to be one of the irreducible polynomials employed in [12].

Example 1: Let $\deg f = 621$, let $f_1 = x^{127} + x^8 + x^7 + x^3 + 1$ and assume $f_1 | f$. If we consider $\hat{e}_i$ to be the i-th entity of the coefficient vector of $\hat{e} = e \bmod f_1$, then we express $\hat{e}_0$ as a sum of bits from e as follows,

$$\begin{aligned}\hat{e}_0 = \quad e_0 + \quad & e_{127} + e_{246} + e_{247} + e_{251} + e_{254} + e_{365} + \\ & e_{367} + e_{375} + e_{381} + e_{484} + e_{485} + e_{486} + \\ & e_{487} + e_{489} + e_{491} + e_{492} + e_{495} + e_{499} + \\ & e_{500} + e_{501} + e_{502} + e_{505} + e_{508} + e_{603} + e_{607}.\end{aligned}$$

In particular, we are interested in linear relations that have as few noise variables from e involved as possible. We use the Piling-up lemma to determine the new bias in $\hat{e}$ after summing up a number of error bits.

Lemma 1 (Piling-up lemma): Let $X_1, X_2, ...X_n$ be i.i.d binary variables where each $\mathbf{Pr}[X_i = 0] = \frac{1}{2}(1 + \epsilon_i)$, for $1 \le i \le n$. Then,

$$\mathbf{Pr}[X_1 + X_2 + \cdots + X_n = 0] = \frac{1}{2}\left(1 + \prod_{i=1}^{n} \epsilon_i\right).$$

[4]The code is a punctured LFSR code.

For instance, the linear relation given in Example 1 has weight 26. Hence, by the Piling-up lemma, the bias of $\hat{e}$ in that particular position (position 0) is ϵ^{26}, i.e., $\mathbf{Pr}[\hat{e}_0 = 1] = \frac{1}{2}(1 - \epsilon^{26})$.

Note: We assume that f_1 is an irreducible polynomial throughout the paper. However, this condition is not a necessity, as the essential feature of the new attack is a CRT map. Actually, it is sufficient if the two polynomials f_1 and f/f_1 are coprime; for example, we could set the polynomial with small degree to be the product of several irreducible polynomials and obtain a solution as well.

B. Using Low-Weight Relations to Build a Distinguisher

We will now show how to build a distinguisher for Ring-LPN with a reducible polynomial using the CRT transformation described in the previous section.

In Example 1, we give a linear relation expressing an error variable $\hat{e}_0$ in the smaller ring as a sum of relatively few error variables in the larger ring. In our example, the polynomial f_1 »behaves well«; it is very sparse and yields a low-weight relation expressing a single noise variable $\hat{e}_0$. However, this will generally not be the case: it may be very difficult to find a desired linear relation with few error variables.

We observe an interesting connection between this problem and searching for codewords with minimum distance in a linear code. The CRT transformation can be viewed as

$$\begin{bmatrix} 1 \bmod f_1 \\ x \bmod f_1 \\ \vdots \\ x^{t-1} \bmod f_1 \end{bmatrix} = \underbrace{\left[\begin{array}{c} \mathbf{g}_0 \\ \mathbf{g}_1 \\ \vdots \\ \mathbf{g}_{l-1} \\ \hline \mathbf{g}_l \\ \mathbf{g}_{l+1} \\ \vdots \\ \mathbf{g}_{t-1} \end{array}\right]}_{=G_\psi^{\mathrm{T}}} \begin{bmatrix} 1 \\ x \\ \vdots \\ x^{l-1} \end{bmatrix},$$

where the top part of G_ψ^{T} is an identity matrix and each row $\mathbf{g}_i$, for $0 \le i \le t-1$, is the coefficient vector of the polynomial $x^i \bmod f_1$.

Thereby, expressing the error polynomial in the smaller ring

$$\begin{aligned} e \bmod f_1 &= \underbrace{\begin{bmatrix} e_0 & e_1 & \cdots e_{t-1} \end{bmatrix}}_{\stackrel{\text{def}}{=}\mathbf{e}^{\mathrm{T}}} G_\psi^{\mathrm{T}} \begin{bmatrix} 1 \\ x \\ \vdots \\ x^{l-1} \end{bmatrix} \\ &= \underbrace{\begin{bmatrix} \hat{e}_0 & \hat{e}_1 & \cdots \hat{e}_{l-1} \end{bmatrix}}_{\stackrel{\text{def}}{=}\hat{\mathbf{e}}^{\mathrm{T}}} \begin{bmatrix} 1 \\ x \\ \vdots \\ x^{l-1} \end{bmatrix}, \end{aligned}$$

we obtain

$$\mathbf{e}^{\mathrm{T}} G_\psi^{\mathrm{T}} = \hat{\mathbf{e}}^{\mathrm{T}} \Rightarrow \hat{\mathbf{e}} = G_\psi \mathbf{e}.$$

Let d_ψ be the minimum distance of the linear code generated by G_ψ. Then by definition, there exists at least one vector $\mathbf{m}$

such that the product $\mathbf{m}G_\psi$ has Hamming weight exactly d_ψ. More specifically,

$$\langle \hat{\mathbf{e}}, \mathbf{m} \rangle = \langle \mathbf{e}, \mathbf{m}G_\psi \rangle ,$$

where $\langle \mathbf{e}, \mathbf{m}G_\psi \rangle$ is a sum of d_ψ noise variables. Thus, according to Piling-up lemma we obtain the following proposition.

Proposition 2 (Estimate of required samples): If the minimum distance of the linear code G_ψ is d_ψ, then the largest bias of some linear combination of noise variables in the smaller ring is no less than ϵ^{d_ψ}.

Consequently, in order to determine the security related to a certain polynomial, we need to determine the minimum distance of the code generated by G_ψ. By applying well-known algorithms such as *information-set decoding* (ISD) algorithms e.g. [24], we can find the minimum distance.

In the example of Lapin, applying ISD algorithms is not necessary since the polynomials f_i, $1 \leq i \leq 5$ are very sparse and admit a very low row-weight of the generator matrix, well below the GV-bound (which is around 154).

C. Recovering the Secret Polynomial

After determining the strongest bias ($\langle \hat{\mathbf{e}}, \mathbf{m} \rangle$), we move to the recovery part.

1) Transforming: Recall that $\deg f_1 = l$. We ask the oracle $\Pi^{\epsilon}_{\text{RING-LPN}}$ for N samples $(\hat{r}_{(i)}, \hat{s} \cdot \hat{r}_{(i)} + \hat{e}_{(i)})$ and then convert each of them to l standard LPN samples by the mapping τ defined in Section II-C1. Write these samples in the matrix form $(\hat{A}_i, \hat{A}_i\hat{\mathbf{s}} + \hat{\mathbf{e}}_i)$. Then, multiplying with vector $\mathbf{m}$, we construct a new LPN sample,

$$(\mathbf{m}\hat{A}_i, \mathbf{m}\hat{A}_i\hat{\mathbf{s}} + \langle \hat{\mathbf{e}}_i, \mathbf{m} \rangle),$$

from each RING-LPN sample. According to Proposition 2, the created samples are with large bias, i.e., no less than ϵ^{d_ψ}.

The overall computational complexity of this step is bounded by,

$$C_1 = Nl^2 + Nl \cdot (l+1) = Nl \cdot (2l+1).$$

2) Birthday: Put these new samples in a data structure that can be accessed in constant time (e.g., a hash table), indexed by its last k entries of the vector $\mathbf{m}\hat{A}_i$. Then a collision between vectors $\mathbf{m}\hat{A}_i$ and $\mathbf{m}\hat{A}_j$ (denoted $\hat{\hat{\mathbf{r}}}_i$ and $\hat{\hat{\mathbf{r}}}_j$, respectively),

$$(\hat{\hat{\mathbf{r}}}_i, \hat{\hat{\mathbf{r}}}_i \cdot \hat{\mathbf{s}} + \langle \hat{\mathbf{e}}_i, \mathbf{m} \rangle) + (\hat{\hat{\mathbf{r}}}_j, \hat{\hat{\mathbf{r}}}_j \cdot \hat{\mathbf{s}} + \langle \hat{\mathbf{e}}_j, \mathbf{m} \rangle) = (\hat{\mathbf{r}}', \hat{\mathbf{r}}' \cdot \hat{\mathbf{s}} + \hat{e}'),$$

yields a vector $\hat{\mathbf{r}}'$ that has at most $l-k$ nonzero positions. The number of such samples is approximately $M = N^2/2^k$. The new samples, such as $(\hat{\mathbf{r}}', \hat{v}')$, depend only on $l-k$ coefficients of the secret $\hat{s}$ and has a bias that is

$$\epsilon' = \epsilon^{2d_\psi}.$$

Calculating the divergence between the distribution of error in $(\hat{\mathbf{r}}', \hat{v}')$ and the uniform distribution, we find that the number of samples required is $M \geq 1/\epsilon^{4d_\psi}$.

Storing the N LPN samples uses lN bit-operations, and performing the birthday procedure requires $lN^2/2^k$ bit-operations. Thus, the total complexity of this step is,

$$C_2 = Nl \cdot (1 + \frac{N}{2^k}).$$

Thus, at this point, we have generated M vector samples $(\hat{\mathbf{r}}'_i, \hat{v}_i)$. All $\hat{\mathbf{r}}'_i$ vectors have dimension no more than $l-k$ as we cancelled out k bits of $\hat{\mathbf{r}}'_i$. Hence, it is enough to consider only $l-k$ bits of $\hat{s}$, i.e., we assume that $\hat{s}$ is of dimension $l-k$. We are then prepared for the final step.

3) Distinguishing the Best Candidate: Group the samples $(\hat{\mathbf{r}}'_i, \hat{v}_i)$ in sets $L(\hat{\mathbf{r}}'_i)$ according to $\hat{\mathbf{r}}'_i$ and then define the function $f_L(\hat{\mathbf{r}}'_i)$ as

$$f_L(\hat{\mathbf{r}}'_i) = \sum_{(\hat{\mathbf{r}}'_i, \hat{v}_i) \in L(\hat{\mathbf{r}}'_i)} (-1)^{\hat{v}_i}.$$

The Walsh transform of f_L is defined as

$$F(\hat{\mathbf{s}}) = \sum_{\hat{\mathbf{r}}'_i} f_L(\hat{\mathbf{r}}'_i)(-1)^{\langle \hat{\mathbf{s}}, \hat{\mathbf{r}}'_i \rangle}.$$

Here we exhaust all the 2^{l-k} candidates of $\mathbf{s}$ by computing the Walsh transform.

Given the candidate $\hat{\mathbf{s}}$, $F(\hat{\mathbf{s}})$ is the difference between the number of predicted 0 and the number of predicted 1 for the bit $\hat{v}'_i + \langle \hat{\mathbf{s}}, \hat{\mathbf{r}}'_i \rangle$. If $\hat{\mathbf{s}}$ is the correct guess, then it is distributed according to $\mathsf{Ber}_{\frac{1}{2}(1-\epsilon^{2d_\psi})}$; otherwise, it is considered random. Thus, the best candidate $\hat{\mathbf{s}}_0$ is the one that maximizes the absolute value of $F(\hat{\mathbf{s}})$, i.e. $\hat{\mathbf{s}}_0 = \arg\max_{\hat{\mathbf{s}} \in F_2^{l-k}} |F(\hat{\mathbf{s}})|$, and we need approximately ϵ^{4d_ψ} samples to distinguish these two cases. Note that false positives are quickly detected in an additional step and this does not significantly increase complexity.

We employ Fast Walsh-Hadamard Transform technique to accelerate the distinguishing step. For well-chosen parameters, the complexity is approximately,

$$C_3 = (l-k)2^{l-k}.$$

Algorithm 1 Partial recovery of RING-LPN

1. (**Preprocessing**) Determine the minimum weight d_ψ of the linear code generated by the CRT transformation and find its corresponding linear relation $\mathbf{m}$.
2. Ask the oracle $\Pi^{\epsilon}_{\text{RING-LPN}}$ for N samples, and then transform each of them to a standard LPN sample with the largest bias, which is no less than ϵ^{d_ψ}.
3. Use the Birthday technique to reduce the secret length at the cost of decreasing the bias.
4. Perform Fast Walsh-Hadamard Transform on the remaining $l-k$ bits of $\hat{\mathbf{s}}$.
5. Output the $\hat{\mathbf{s}}_0$ that maximizes the absolute value of the transform.

From the new algorithm, there are some important consequences to consider when choosing parameters to thwart our attack. We give some very brief comments, assuming that every smaller ring is of approximately the same size:

1) Choosing a large number of factors in f seems a bit dangerous, as the dimension of the code G_ψ becomes small. In our attack we used a birthday argument to reduce the dimension of $\hat{s}$, but this might not be necessary if the dimension is already very low. Instead we may search for special $\hat{r}$ (e.g. many $\hat{r} = 1$) values that allows quick recovery of $\hat{s}$.

2) One should use irreducible polynomials f_i with degree around $\frac{n}{m}$ for $1 \leq i \leq m$ such that for every f_i the corresponding linear code G_ψ has minimum distance as large as possible. From the GV bound, we know roughly what to expect. However, the GV-bound may not be asymptotically true for codes generated by the CRT-transform.
3) Following this line, a necessary but probably insufficient condition on ϵ is that $1/\epsilon^{4d_\psi} \geq 2^b$ for b-bit security[5].

IV. The Improved Version for the Proposed Instance of Lapin

In [12], Heyse et al. employ the following reducible polynomial,

$$f = \underbrace{\left(x^{127}+x^8+x^7+x^3+1\right)}_{f_1} \cdot \underbrace{\left(x^{126}+x^9+x^6+x^5+1\right)}_{f_2} \cdot \underbrace{\left(x^{125}+x^9+x^7+x^4+1\right)}_{f_3} \cdot \underbrace{\left(x^{122}+x^7+x^4+x^3+1\right)}_{f_4} \cdot \underbrace{\left(x^{121}+x^8+x^5+x^1+1\right)}_{f_5},$$

as the underlying structure of the quotient ring R. This parameter setting is then adopted in a more recent paper [9] by Gaspar at al., to show that the original protocol and its hardware variant, Mask-Lapin, have much gain, compared with the implementation from block ciphers (e.g., AES), in the sense of resisting power analysis attacks by masking.

However, in the sense of thwarting the new attack of this paper, it is not a good selection. We give two reasons as follows.

- As stated previously, for any polynomial f_i ($1 \leq i \leq 5$), the generator matrix G_ψ of the corresponding code is sparse, hence yielding that we could roughly adopt one row vector in G_ψ as its minimum weight codeword. Then, the vector $\mathbf{m}$ is of Hamming weight 1 and we could save the computational cost for linearly combining several samples to form a new one with the largest bias.
- Secondly, for the Lapin instance, the largest bias always holds at the last row of the generator matrix G_ψ, when modulo operation is taken over each irreducible factor. Furthermore, for RING-LPN samples $(\hat{r}_{(i)}, \hat{s}\cdot\hat{r}_{(i)}+\hat{e}_{(i)})$, if we find collisions on the last k positions of $\hat{r}_{(i)}$ (k is larger than 10), then the last row vector in the matrix form of the merged sample $(\hat{r}', \hat{r}'\cdot\hat{s}+\hat{e}')$ is of the form

$$\begin{bmatrix} 0 & \dots & 0 & \hat{r}'_{l-k-1} & \hat{r}'_{l-k-2} & \cdots & \hat{r}'_0 \end{bmatrix}. \quad (2)$$

  This vector can be read from the polynomial $\hat{r}'$ directly, without any computation cost.

Therefore, we could present a specific algorithm for solving the Lapin instance, see Algorithm 2, which is more efficient than the generic one. After determining the weight w of the last row vector in the generator matrix G_ψ, we ask the oracle $\Pi^{\epsilon}_{\text{RING-LPN}}$ for N samples $(\hat{r}_{(i)}, \hat{s}\cdot\hat{r}_{(i)}+\hat{e}_{(i)})$ and search for collisions by the last k coefficients of $\hat{r}$ directly. Then, for each collision represented by a merged RING-LPN sample $(\hat{r}', \hat{v}')$, we construct a new standard LPN sample, where the vector is generated by (2), and the observed value is the coefficient of x^{l-1} in the polynomial $\hat{v}'$. These samples are with the largest bias. The distinguishing step is the same as that in the generic algorithm and we present the detailed complexity analysis and numerical results in the consecutive sections.

[5]This is just one concern as there may be many other aspects to be taken into consideration that will make the problem solvable. For instance, if d_ψ and ϵ are very large the constraint can be satisfied while the problem in a larger ring remains easy.

Algorithm 2 Improved partial key recovery for Lapin

1 (**Preprocessing**) Find the weight w of the $(l-1)$-th row in the generator matrix of the CRT transformation.
2 Ask the oracle $\Pi^{\epsilon}_{\text{RING-LPN}}$ for N samples, index them by the last k coefficients of $\hat{r}$, and then search for all collisions $(\hat{r}', \hat{r}'\cdot\hat{s}+\hat{e}_1+\hat{e}_2) = (\hat{r}_1, \hat{s}\cdot\hat{r}_1+\hat{e}_1) + (\hat{r}_2, \hat{s}\cdot\hat{r}_2+\hat{e}_2)$, where the last k coefficients of $\hat{r}_1$ and $\hat{r}_2$ are the same.
3 Generate standard LPN samples from the RING-LPN samples, and then perform Fast Walsh-Hadamard Transform on $l-k$ bits of $\hat{\mathbf{s}}$.
4 Output the $\hat{\mathbf{s}}_0$ that maximizes the absolute value of the transform.

V. Complexity Analysis

Having introduced the two versions in detail, we now estimate their complexity. It is straightforward that the generic algorithm (Algorithm 1) costs $C = C_1 + C_2 + C_3$ bit-operations. But analyzing the improved attacking complexity for the proposed instance in Lapin is more attractive, which is stated in the following theorem.

Theorem 2 (The complexity of Algorithm 2): Let w be the weight of the last row in the CRT transformation matrix G_ψ. Then, the complexity of Algorithm, denoted C^*, is given by

$$C^* = l\cdot(N + N^2/2^k) + (l-k)2^{l-k}, \quad (3)$$

under the condition that $N^2/2^k \geq \frac{1}{\epsilon^{4w}}$.

Proof: We analyze step by step: 1) First, we store the samples received from N oracle calls into a table using lN bit-operations. 2) Clearing k bits in a collision procedure yields $N^2/2^k$ samples and can be performed in $lN^2/2^k$ bit-operations. As the required standard LPN instance can be read directly from the RING-LPN instance, the transformation has no computational cost. 3) Afterwards, we perform a Fast Walsh-Hadamard Transform. If the number of unknown bits $l-k$ is at least $\log_2\left(N^2/2^k\right)$, then the complexity is $(l-k)2^{l-k}$. This is the final step. Summarizing the individual steps yields C^* which finalizes the proof. ■

VI. Results

We now present numerical results of the improved partial key recovery attack on the authentication protocol Lapin [12]. The attack we describe concerns the instance of Lapin using the degree 621 polynomial f and with the parameter $\eta = \frac{1}{6}$. As claimed in [12], this given instance is designed to resist the best known attack on RING-LPN within the complexity 2^{80}. However, we have shown that it is possible to greatly improve the attack complexity. The improvements can be seen in Table II.

TABLE II
THE COMPLEXITY FOR ATTACKING EACH MODULUS OF THE PROPOSED INSTANCE USING REDUCIBLE POLYNOMIAL f.

Polynomial f_i	Parameters			$\log_2 C^*$
	k	w	$\log_2 N$	
$x^{127}+x^8+x^7+x^3+1$	65	26	63	70.56
$x^{126}+x^9+x^6+x^5+1$	63	26	62	70.30
$x^{125}+x^9+x^7+x^4+1$	63	26	62	69.96
$x^{122}+x^7+x^4+x^3+1$	60	27	62	75.02
$x^{121}+x^8+x^5+x^1+1$	58	29	63	71.31

For the distinguishing attack, the complexity is only $2^{69.96}$. For actual recovery of the secret polynomial, we need all five coordinates in the CRT representation, which gives an upper bound on the security that is roughly $2^{75.05}$. However, to solve this search problem, we describe a specified approach presented in the appendix that reduces the complexity to about $2^{71.88}$. A comparison with previous algorithms is given in Table I in the introduction.

VII. CONCLUDING REMARKS

We have proposed a new generic algorithm to solve the reducible case of RING-LPN. By exploiting the ring structure further, our new algorithm is much more efficient than previous algorithms, enough to break the claimed 80-bit security for one of the two proposed instances of Lapin.

We have shown that a linear code arising from the CRT transform characterizes the performance of this attack through its minimum distance. This is combined with some standard techniques of using birthday or possibly generalized birthday arguments and efficient recovery through Fast Walsh-Hadamard transform.

The low-weight property of the polynomials in the Lapin case makes the problem considerably easier than otherwise and thus makes Lapin susceptible to our attack. Using really low-weight irreducible polynomials such as $x^{127}+x+1$ can give rise to linear relations with weight as low as 10 or even less. We have not seen that such polynomials have been pointed out as very weak before.

The description of the new algorithm was influenced by the Lapin case. There are more improvements that can be described in the general case. One such improvement is the use of a generalized birthday technique [26]. This will allow us to consider larger dimensions at the cost of increasing the noise level. We have also noted that the simple bit-oriented samples in this paper can be replaced by more complicated vectorial samples, which will give a stronger bias.

APPENDIX A
LAPIN AUTHENTICATION PROTOCOL

In this section, we describe the Lapin two-round authentication protocol. Let $R=\mathbb{F}_2[x]/(f)$ be a ring and R^* the set of units in R. The protocol is defined over the ring R. Let π be a mapping, chosen such that for all $c,c'\in\{0,1\}^\lambda$, $\pi(c)-\pi(c')\in R\backslash R^*$ if and only if $c=c'$. Furthermore, let $\eta\in\left(0,\frac{1}{2}\right)$ be a Bernoulli distribution parameter and $\eta'\in\left(\eta,\frac{1}{2}\right)$ a threshold parameter. The elements R, $\pi:\{0,1\}^\lambda\to R$, η and η' are public parameters. The ring elements $s,s'\in R$ constitute the secret key.

Protocol 1 Lapin two-round authentication

Tag		**Reader**
		$c\xleftarrow{\$}\{0,1\}^\lambda$
	$\xleftarrow{c}$	
$r\xleftarrow{\$}R^*; e\xleftarrow{\$}\mathsf{Ber}_\eta^R$		
$z\leftarrow r\cdot(s\cdot\pi(c)+s')+e$		
	$\xrightarrow{(r,z)}$	
		if $r\notin R^*$ then reject
		$e'\leftarrow z-r\cdot(s\cdot\pi(c)+s')$
		if $w_{\mathrm{H}}(e')>n\cdot\eta'$ then reject
		else accept

Suppose that we have a key-generation oracle; it will give us the key $s,s'\xleftarrow{\$}R$. The secret key is shared among the tag and the reader. Protocol 1 gives how information is exchanged between the tag and the reader in Lapin.

APPENDIX B
A SPECIAL SECRET RECOVERY APPROACH

In this section, we describe a better secret recovering approach on the instance proposed in Lapin with a reducible polynomial. This attack exploits the secret polynomial's representations in the three relatively easy-attacked quotient rings $\mathbb{F}_2[x]/(f_i), i=1,2,3$, to decrypt those in rings $\mathbb{F}_2[x]/(f_i), i=4,5$, and thus obtains higher efficiency than simply attacking one ring by another. Actually, it reduces the attacking complexity from $2^{75.05}$ to $2^{71.88}$.

Denote $\mathbf{a}||\mathbf{b}$ as the concatenation of two vectors $\mathbf{a}$ and $\mathbf{b}$, and $(s_1 \bmod f_1, s_2 \bmod f_2,\ldots,s_5 \bmod f_5)$ as the CRT representation of s. The three-step attack is described as follows.

We first recover the secret polynomials s_1, s_2 and s_3 in the corresponding smaller rings. This step costs around $2^{71.88}$ bit operations. Then, using these known polynomials, we show that the secrets s_4 and s_5 can be recovered with negligible costs.

The second step is transforming each RING-LPN sample to its corresponding standard LPN samples. This step will be tricky as we want to make use of the known information. Since f_i are distinct irreducible polynomials over $\mathbb{F}_2[x]$, there exist polynomials $t_i\in R$ such that $s=\sum_{i=1}^5 s_i\cdot t_i$ and $t_i\equiv 1 \bmod f_i$, where $i=1,2,\ldots,5$. Moreover, these polynomials $t_i, i=1,2,\ldots,5$ can be computed efficiently. Then, for each

sample (r, v), where $v = r \cdot s + e$, by adding the polynomial $r \cdot (\sum_{i=1}^{3} s_i \cdot t_i)$ to v, we create a new polynomial $r \cdot (s_4 \cdot t_4 + s_5 \cdot t_5) + e$, which will be converted to a vector $A_4\mathbf{s}_4 + A_5\mathbf{s}_5 + \mathbf{e}$, where $A_4(A_5)$ is the matrix whose i-th column is the coefficient vector of $r \cdot t_4 \cdot x^i (r \cdot t_5 \cdot x^i) \bmod f$. The degree of $s_4(s_5)$ is 122(121), thereby yielding that only the first 122(121) columns of the matrix $A_4(A_5)$, denoted $A'_4(A'_5)$, are useful. We can compute those sub-matrices in $2^{17.2}$ bit operations for each RING-LPN sample, thereby economically constructing 621 standard LPN samples $(\mathbf{r}_4^i || \mathbf{r}_5^i, \langle \mathbf{s}_4 || \mathbf{s}_5, \mathbf{r}_4^i || \mathbf{r}_5^i \rangle + e_i)$ for $i = 0, 1, \ldots, 620$, where $\mathbf{r}_4^i(\mathbf{r}_5^i)$ is the i-th row in the matrix $A'_4(A'_5)$.

The last step is to attack a standard LPN problem with length 243 and error probability $1/6$. By Levieil-Fouque algorithm, the cost is small, i.e., around $2^{57.1}$ bit operations. It requires $2^{46.6}$ standard LPN samples, i.e. $2^{37.3}$ RING-LPN ones.

The number of samples in the second step is also bounded by that required in the final step; the overall cost of the last two steps, therefore, is no more than 2^{60} bit operations, far less than that required in the starting step. Thus, it is reasonable to embrace $2^{71.88}$ bit operations as the total attacking complexity.

REFERENCES

[1] Blum, A., Furst, M., Kearns, M., Lipton, R.: Cryptographic Primitives Based on Hard Learning Problems. In: Stinson, R. (ed.) CRYPTO 1993, LNCS, vol. 773, pp. 278–291. Springer, Heidelberg (1994)

[2] Blum, A., Kalai, A., Wasserman, H.: Noise-Tolerant Learning, the Parity Problem, and the Statistical Query Model. In: Journal of the ACM, vol. 50, no. 4, pp. 506–519. (2003)

[3] Berlekamp, E.R., McEliece, R.J., van Tilborg, H.C.A.: On the Inherent Intractability of Certain Coding Problems. IEEE Trans. Info. Theory, vol. 24, pp. 384–386. (1978)

[4] Bernstein, D., Lange T.: Never trust a bunny. In: Radio Frequency Identification Security and Privacy Issues, pp. 137–148. Springer, Berlin Heidelberg (2013)

[5] Chose, P., Joux, A., Mitton, M.: Fast Correlation Attacks: An Algorithmic Point of View. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 209–221. Springer, Heidelberg (2002)

[6] Damgard, I., Park, S.: Is Public-Key Encryption Based on LPN Practical? Cryptology ePrint Archive, Report 2012/699 (2012), http://eprint.iacr.org/

[7] Dodis, Y., Kiltz, E., Pietrzak, K., Wichs, D.: Message Authentication, Revisited. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 355–374. Springer, Heidelberg (2012)

[8] Fossorier, M.P.C., Mihaljevic, M.J., Imai, H., Cui, Y., Matsuura, K.: A Novel Algorithm for Solving the LPN Problem and its Application to Security Evaluation of the HB Protocol for RFID Authentication. Cryptology ePrint archive, Report 2012/197 (2012), http://eprint.iacr.org/

[9] Gaspar, L., Leurent, G., Standaert, F. X.: Hardware Implementation and Side-Channel Analysis of Lapin. CT-RSA 2014, (2014)

[10] Gilbert, H., Robshaw, M.J.B., Seurin, Y.: HB$^{\#}$: Increasing the Security and the Efficiency of HB^{+}. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 361–378. Springer, Heidelberg (2008)

[11] Gilbert, H., Robshaw, M.J.B., Sibert, H.: An active attack against HB^{+}—a provably secure lightweight authentication protocol. Cryptology ePrint Archive, Report 2005/237 (2005), http://eprint.iacr.org/

[12] Heyse, S., Kiltz, E., Lyubashevsky, V., Paar, C., Pietrzak, K.: Lapin: An Efficient Authentication Protocol Based on Ring-LPN. In: FSE 2012, pp. 346–365. (2012)

[13] Hopper, N.J., Blum, M.: Secure human identification protocols. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 52–66. Springer, Heidelberg (2001)

[14] Juels, A., Weis, S.A.: Authenticating pervasive devices with human protocols. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 293–308. Springer, Heidelberg (2005)

[15] Katz, J., Shin, J.S.: Parallel and concurrent security of the HB and HB^{+} protocols. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 73–87. Springer, Heidelberg (2006)

[16] Katz, J., Shin, J.S., Smith, A.: Parallel and concurrent security of the HB and HB^{+} protocols. Journal of Cryptology 23(3), 402–421 (2010)

[17] Kearns., M: Effcient Noise-Tolerant Learning from Statistical Queries. In: J. ACM 45(6), pp. 983–1006. (1998)

[18] Kiltz, E., Pietrzak, K., Cash, D., Jain, A., Venturi, D.: Efficient Authentication from Hard Learning Problems. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 7–26. Springer, Heidelberg (2011)

[19] Kirchner, P.: Improved Generalized Birthday Attack. Cryptology ePrint Archive, Report 2011/377 (2011), http://eprint.iacr.org/

[20] Levieil, E., Fouque, P. A.: An Improved LPN Algorithm. In: Proceedings of SCN 2006, LNCS 4116, pp. 348–359. Springer, Heidelberg (2006)

[21] Lyubashevsky, V.: The Parity Problem in the Presence of Noise, Decoding Random Linear Codes, and the Subset Sum Problem, In: Chekuri, C., Jansen, K., Rolim, J.D.P., Trevisan, L. (eds.) APPROX-RANDOM 2005, LNCS, vol. 3624, pp. 378–389. Springer, Heidelberg (2005)

[22] Munilla, J., Peinado, A.: HB-MP: A further step in the HB-family of lightweight authentication protocols. Computer Networks 51(9), pp. 2262–2267. (2007)

[23] Regev, O.: On Lattices, Learning with Errors, Random Linear Codes, and Cryptography. In: Gabow, H.N., Fagin, R. (eds.) 37th Annual ACM Symposium on Theory of Computing, Proceedings, pp. 84–93. (2005)

[24] Stern, J.: A Method for Finding Codewords of Small Weight. In: Wolfmann, J., Cohen, G. (eds.) Coding Theory 1988. LNCS, vol. 388, pp. 106–113. Springer, Heidelberg (1989)

[25] Stern, J.: A New Identification Scheme Based on Syndrome Decoding. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 13–21. Springer, Heidelberg (1994)

[26] Wagner, D.: A Generalized Birthday Problem. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 288–304. Springer, Heidelberg (2002)